

a

1. (currently amended) In a database system, a method of maintaining a self-tuning histogram having a plurality of existing buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range comprising ~~the step of~~ creating at least one new bucket in response to a query on the database wherein the new bucket has bucket boundaries corresponding to a range of tuple attribute values returned by the query and a bucket frequency corresponding to a number of tuples returned by the query and wherein each new bucket is contained within at least one existing bucket and wherein the new bucket becomes a child bucket that has a child bucket frequency and the existing bucket becomes a parent bucket that has a parent bucket frequency.

2. (original) The method of claim 1 wherein each bucket has a rectangular shape.

3. (canceled)

4. (original) The method of claim 1 wherein a total number of buckets is limited to a predetermined budget.

5. (currently amended) The method of claim 4 comprising ~~the step of~~ merging buckets based on a merge criterion when the total number of buckets exceeds the predetermined budget.

6. (original) The method of claim 5 wherein the merge criterion is a similar bucket density, wherein bucket density is based on the bucket frequency divided by the bucket volume.

7. (currently amended) The method of claim 1 further comprising ~~the step of~~ shrinking the boundaries of each new bucket if the boundaries of the new bucket intersect any existing bucket boundaries.

8. (original) The method of claim 1 wherein the frequency of the parent bucket is diminished by the frequency of the child bucket.

9. (currently amended) In a database system, a method of maintaining a self-tuning histogram having a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attributes that fall in the bucket boundary range comprising the steps of:

- a) examining the results of a query executed on the database;
- b) creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query;
- c) modifying the boundaries of each candidate hole such that the boundaries of the modified hole is are completely contained within the boundaries of at least one existing parent bucket and does do not partially intersect the boundaries of any existing bucket; and
- d) creating a new child bucket that has a child frequency in the histogram corresponding to each modified hole.

10. (original) The method of claim 9 wherein each bucket has a rectangular shape.

11. (canceled)
12. (original) The method of claim 9 wherein a total number of buckets is limited to a predetermined budget.

13. (currently amended) The method of claim 12 comprising ~~the step of~~ merging buckets based on a merge criterion when the total number of buckets exceeds the predetermined budget.

14. (original) The method of claim 13 wherein the merge criterion is a similar bucket density, wherein bucket density is based on the bucket frequency divided by the bucket volume.

15. (original) The method of claim 9 wherein the frequency of the parent bucket is diminished by the frequency of the child bucket.

16. (currently amended) A computer readable medium having computer executable instructions for performing steps for maintaining a self-tuning histogram having a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range, the steps comprising:

- a) examining the results of a query executed on the database;
- b) creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query;
- c) modifying the boundaries of each candidate hole such that the boundaries of the modified hole is are completely contained within the boundaries of at least one existing parent bucket and does do not partially intersect the boundaries of any existing bucket; and

d) creating a new child bucket that has a child frequency in the histogram corresponding to each modified hole.

17. (original) The computer readable medium of claim 16 wherein each bucket has a rectangular shape.

18. (canceled)

19. (currently amended) The computer readable medium of claim 16 wherein the method steps comprise comprising the step of merging buckets having a similar bucket density when the total number of buckets exceeds a predetermined budget.

20. (currently amended) An apparatus for maintaining a self-tuning histogram having a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range comprising:

a) means for examining the results of a query executed on the database;

b) means for creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query;

c) means for modifying the boundaries of each candidate hole such that the boundaries of the modified hole is are completely contained within the boundaries of at least one existing parent bucket and does do not partially intersect the boundaries of any existing bucket; and

d) means for creating a new child bucket that has a child frequency in the histogram corresponding to each modified hole.

21. (canceled)

22. (currently amended) An apparatus that maintains ~~for maintaining~~ a self-tuning histogram having a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range comprising:

a) a memory device that stores ~~for storing~~ a database comprising multiple data records;

b) a computer having one or more processing units that execute ~~for executing~~ a stored computer program, said computer including a rapid access memory store; and

c) an interface that couples ~~for coupling~~ the memory device that stores ~~for storing~~ the database to the computer to allow records to be retrieved from the database; wherein

d) the stored program has components including i) a component ~~for examining~~ that examines the results of a query executed on the database; ii) a component that creates ~~for creating~~ at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query; iii) a component ~~for modifying~~ that modifies the boundaries of each candidate hole such that the boundaries of the modified hole is ~~are~~ completely

contained within the boundaries of at least one existing parent bucket and does do not partially intersect the boundaries of any existing bucket; and iv) a component that creates ~~for creating~~ a new child bucket that has a child frequency in the histogram corresponding to each modified hole.

23. (canceled)
24. (new) For use with a database system, a histogram tuning system comprising:
- a component that receives a bucket from the histogram; and
- a tuning component that iteratively populates the bucket with a child bucket, as a function of query results, wherein the child bucket is completely contained within the bucket.
25. (new) The histogram tuning system of claim 24 wherein the tuning component populates the bucket with a child bucket that has boundaries corresponding to a range of attribute values present in the query results and a child bucket frequency corresponding to a number of tuples present in the query results.
26. (new) The histogram tuning system of claim 24 comprising a merging component that merges buckets based on a merge criteria.
27. (new) The histogram tuning system of claim 26 wherein the merge criterion is a similar bucket density.
28. (new) The histogram tuning system of claim 24 wherein the tuning component shrinks the boundaries of the child bucket if the child bucket boundaries intersect any other bucket boundaries.

29. (new) The histogram tuning system of claim 25 wherein the bucket has a frequency and the bucket frequency is diminished by the child bucket frequency.
30. (new) A database histogram tuning system comprising:
means for receiving a bucket from a histogram; and
means for iteratively populating the bucket with a child bucket, as a function of query results, such that the child bucket is fully contained within the received bucket.
-